FAST ALGORITHM FOR LOW-RANK MATRIX RECOVERY IN POISSON NOISE

Yang Cao and Yao Xie

Industrial and System Engineering, Georgia Institute of Technology

arXiv:1407.0726v2 [stat.ML] 19 Dec 2014

ABSTRACT

This paper describes a new algorithm for recovering low-rank matrices from their linear measurements contaminated with Poisson noise: the Poisson noise Maximum Likelihood Singular Value thresholding (PMLSV) algorithm. We propose a convex optimization formulation with a cost function consisting of the sum of a likelihood function and a regularization function which is proportional to the nuclear norm of the matrix. Instead of solving the optimization problem directly by semi-definite program (SDP), we derive an iterative singular value thresholding algorithm by expanding the likelihood function. We demonstrate the good performance of the proposed algorithm on recovery of solar flare images with Poisson noise: the algorithm is more efficient than solving SDP using the interior-point algorithm and it generates a good approximate solution compared to that solved from SDP.

Index Terms— low-rank matrix recovery, nuclear norm, singular value thresholding, solar flare images

1. INTRODUCTION

Recovery of a matrix M from its linear measurements (or linear projections) contaminated with Poisson noise arises from various important applications such as optical imaging, nuclear medicine and X-ray imaging [1]. When M is low-rank, we can still recover M from a relatively small number of measurements, and it has been shown that under certain conditions M can be recovered exactly [2].

While there has been much success for low-rank matrix recovery and completion without noise or with additive Gaussian noise, relatively fewer results are available when the measurements are contaminated with Poisson noise [3, 4]. In [3], the authors establish some performance bounds without developing algorithms. In [4], the parameters of Poisson distribution is sparse coefficients under certain compression dictionary, which is different with that of our model. The problem with Poisson noise is different because unlike Gaussian noise which has static noise variance, the variance of Poisson noise is proportional to the signal intensity. Also, we need to use a non-linear likelihood function to replace the ℓ_2 norm penalty for data fitting term in the formulation. Moreover, in practical systems with Poisson noise, many physical constraints have to be taken into consideration when recover the signal, e.g. the positivity of the signal and the total intensity constraint.

In this paper, we present a regularized maximum likelihood estimator to recover a low-rank matrix from linear measurements contaminated with Poisson noise. Instead of directly solving the convex optimization problem formulated this way, we present a generalized iterative singular value thresholding method [5], which can be viewed as a consequence of approximating the log likelihood function by its second order Taylor expansion. The good performance of the proposed algorithm is demonstrated via numerical examples where we recover solar flare images with low-rank structure from Poisson measurements. We show that the proposed method is more efficient than solving the convex optimization using interior point method and it has good accuracy.

Submitted to IEEE GLOBALSIP 2014, May 2014.

2. FORMULATION

2.1. Model

Suppose we wish to recover a matrix $M^* \in \mathbb{R}_+^{m_1 \times m_2}$ consisting of nonnegative entries from N linear measurements with Poisson measurements $y_i \in \mathbb{Z}_+^N$ that take the forms of

$$y_i \sim \text{Poisson}([\mathcal{A}M^*]_i), \quad i = 1, \ldots, N, \tag{1}$$

where the linear operator $\mathcal{A} : \mathbb{R}_+^{m_1 \times m_2} \to \mathbb{R}^N$ models the measuring process of physical devices and it takes the following form:

$$[\mathcal{A}M]_i = \langle A_i, M \rangle \triangleq \text{tr}(A_i^\top M), \tag{2}$$

where $A_i \in \mathbb{R}^{m_1 \times m_2}$, and $\text{tr}(X)$ denotes trace of a matrix X. In optical systems, the matrix A_i models the masks that are applied to the light field before the intensity is measured. Let $\text{vec}(X) = [x_1^\top, \ldots, x_n^\top]^\top$ denote vectorized version of the matrix $X = [x_1, \ldots, x_n]$. Note that if we define

$$A \triangleq \begin{bmatrix} \text{vec}(A_1)^\top \\ \vdots \\ \text{vec}(A_N)^\top \end{bmatrix}, \quad f \triangleq \text{vec}(M),$$

then the measurements can be written as

$$\mathcal{A}M = Af.$$

We make the following assumptions about the system. Let $[X]_{ij}$ denote the element of matrix X in the ith row and jth column, and $[x]_j$ denotes the jth element of a vector x. Define the norm $\|X\|_{1,1} = \sum_i \sum_j [X]_{ij}$. By this notation, we assume that the total intensity of M^*, given by

$$I \triangleq \|M^*\|_{1,1} \tag{3}$$

is known a priori. Also, to have physically realizable linear optical systems, we assume that the measurement operator $\mathcal{A}$ satisfies the following constraints:

1. (positivity-preserving) $[M]_{ij} \geq 0$ for all $i, j \Rightarrow [\mathcal{A}M]_i \geq 0$, for all i,
2. (flux-preserving) $\sum_{i=1}^N [\mathcal{A}M]_i \leq \|M\|_{1,1}$.

Our goal is to estimate the signal $M^* \in \mathbb{R}_+^{m_1 \times m_2}$ from measurements $y \in \mathbb{Z}_+^N$.

2.2. Regularized Maximum-Likelihood Estimator

We propose a regularized maximum-likelihood estimator. The probability density function of y is given by

$$p(y|\mathcal{A}M^*) = \prod_{j=1}^{N} \frac{[\mathcal{A}M^*]_j^{y_j}}{y_j!} e^{-[\mathcal{A}M^*]_j}. \quad (4)$$

The corresponding regularization likelihood function is given by the following optimization problem

$$\begin{aligned}\widehat{M} &\triangleq \arg\min_{M\in\Gamma}[-\log p(y|\mathcal{A}M) + \lambda\rho(M)] \\ &= \arg\min_{M\in\Gamma}[-\sum_{j=1}^{N} y_j \log[\mathcal{A}M]_j + [\mathcal{A}M]_j + \lambda\rho(M)], \quad (5)\end{aligned}$$

where $\rho(M) > 0$ is a regularization function and $\lambda > 0$ is the regularization parameter. Here Γ is a countable set of feasible estimators

$$\Gamma \triangleq \{M_i \in \mathbb{R}_+^{m_1\times m_2} : \|M_i\|_{1,1} = I, i = 1, 2, \ldots\}, \quad (6)$$

and the regularization function satisfies the Kraft inequality

$$\sum_{M\in\Gamma} e^{-\rho(M)} \le 1. \quad (7)$$

We can think of this formulation as a discretized feasible domain version of the general regularized maximum likelihood estimator. The regularization function assigns small value for lower rank M and vice versa. Using Kraft-compliant regularization to prefix codes for estimators is a commonly used technique [6].

The performance metric we use for estimator is a normalized risk:

$$R(M^*, M) \triangleq \frac{1}{I^2}\|M^* - M\|_F^2, \quad (8)$$

where $\|X\|_F^2$ is the Frobenius norm of a matrix X, which is defined as $\|X\|_F^2 = \sum_i \sum_j |[X]_{ij}|^2$.

2.3. Sensing operator

We adopt the following linear sensing operator $\mathcal{A}$ by using the sensing matrices suggested by [6] with the form of

$$[A_i]_{jk} = \begin{cases} 0, & \text{with probability } p; \\ 1/N, & \text{with probability} 1-p. \end{cases} \quad (9)$$

It can be verified that this operator $\mathcal{A}$ satisfies the requirements in the previous subsection. In particular,
1. all entries of A_i take values of 0 or $1/N$;
2. $\mathcal{A}$ satisfies flux preserving. Since all entries of A_i are less than $1/N$, for $[M]_{ij} \ge 0$ for all i, j,

$$\|\mathcal{A}M\|_1 = \sum_{i=1}^{N}\sum_{j=1}^{m_1}\sum_{k=1}^{m_2}[A_i]_{jk}M_{jk} \le \sum_{j=1}^{m_1}\sum_{k=1}^{m_2} M_{jk} = I.$$

where $\|x\|_1 = \sum_i |x_i|$ denotes the ℓ_1 norm of a vector x.
3. with probability at least $1 - Np^{m_1+m_2}$, every A_i has at least one non-zero entry. It follows that for M such that $[M]_{ij} \ge c$, for all i, j, we have that

$$[\mathcal{A}M]_i = \sum_{j=1}^{m_1}\sum_{k=1}^{m_2}[A_i]_{jk}M_{jk} \ge c\sum_{j=1}^{m_1}\sum_{k=1}^{m_2}[A_i]_{jk} \ge c/N.$$

This operator also satisfies the restrictive isometry property.

3. PMLSV ALGORITHM

In this section, we introduce a Poisson noise Maximal Likelihood Singular Value thresholding (PMLSV) algorithm for solving the regularized maximum likelihood problem formulated in (5).

Nuclear norm of some matrix is defined as sum of singular values of the matrix and it is proven to be a very useful norm when solving low-rank matrix recovery problem because its close connection with rank of matrix [2] and its convexity. Therefore, it is reasonable for us to use the nuclear norm of M, denoted as $\|M\|_*$, for $\rho(M)$. Therefore, we recover the low-rank matrix by solving an optimization problem

$$\min_{M\in\Gamma_0} f(M) + \lambda\|M\|_* \quad (10)$$

where $f(M) = -\log p(y|\mathcal{A}M)$, and we also relax the feasible domain Γ to Γ_0

$$\Gamma_0 \triangleq \{M \in \mathbb{R}_+^{m_1\times m_2} : \|M\|_{1,1} = I\}. \quad (11)$$

To solve (10), we may use the interior-point method since nuclear norm minimization problem with convex feasible domain can be reformulated as a Semidefinite program (SDP). However, the large number of dummy variables makes this approach less preferable as an computationally efficient algorithm for large problem. Hence, we seek an alternative approximate algorithm other than solving the SDP.

To derive the approximate algorithm, we first expand the likelihood function of our cost function by Taylor expansion and only keep up to second term as approximation. Under such approximation, (10) becomes

$$M_k = \arg\min_M \left[Q_{t_k}(M, M_{k-1}) + \lambda\|M\|_*\right], \quad (12)$$

with

$$\begin{aligned}Q_{t_k}(M, M_{k-1}) &:= f(M_{k-1}) + \langle M - M_{k-1}, \nabla f(M_{k-1})\rangle \\ &\quad + \frac{t_k}{2}\|M - M_{k-1}\|_F^2, \quad (13)\end{aligned}$$

where t_k is the step size at kth iteration. Note that (13) is an optimization problem with a form similar to that studied in [8] and the optimizer can be derived analytically as follows. By dropping and introducing terms independent on M whenever needed (denoted by "$\propto$"), we can rewrite $Q_{t_k}(M, M_{k-1})$ as:

$$\begin{aligned}&Q_{t_k}(M, M_{k-1}) \\ &\propto \langle M - M_{k-1}, \nabla f(M_{k-1})\rangle + \frac{t_k}{2}\|M - M_{k-1}\|_F^2 \\ &\propto \langle M - M_{k-1}, \nabla f(M_{k-1})\rangle + \frac{t_k}{2}\|M - M_{k-1}\|_F^2 \\ &\quad + \frac{t_k}{2}\|\frac{1}{t_k}\nabla f(M_{k-1})\|_F^2 \\ &\propto \frac{t_k}{2}\left\langle M - M_{k-1} + \frac{1}{t_k}\nabla f(M_{k-1}), M - M_{k-1} + \frac{1}{t_k}\nabla f(M_{k-1})\right\rangle \\ &= \frac{t_k}{2}\left\|M - \left(M_{k-1} - \frac{1}{t_k}\nabla f(M_{k-1})\right)\right\|_F^2 \quad (14)\end{aligned}$$

Substituting (14) into (12) and scale the cost function by $1/t_k$, we

have:

$$M_k = \arg\min_M \left[\frac{1}{2} \left\| M - \left(M_{k-1} - \frac{1}{t_k} \nabla f(M_{k-1}) \right) \right\|_F^2 + \frac{\lambda}{t_k} \|M\|_* \right]. \quad (15)$$

The solution to (15) is given by a form of Singular Value Thresholding (SVT) [5]. Consider the following problem

$$\min_{Y \in \mathbb{R}^{n_1 \times n_2}} \left\{ \frac{1}{2} \|Y - X\|_F^2 + \tau \|Y\|_* \right\}, \quad (16)$$

where $X \in \mathbb{R}^{n_1 \times n_2}$ is given and τ is the regularization parameter. For a matrix $X \in \mathbb{R}^{n_1 \times n_2}$ with rank r, let its singular value decomposition be $X = U\Sigma V^T$, where $U \in \mathbb{R}^{n_1 \times r}$, $V \in \mathbb{R}^{n_2 \times r}$, $\Sigma = \text{diag}(\{\sigma_i\}, i = 1, 2, ..., r)$, and σ_i is a singular value of the matrix X. For each $\tau \geq 0$, define the singular value thresholding operator as :

$$D_\tau(X) \triangleq U D_\tau(\Sigma) V^T, \quad (17)$$

where the $D_\tau(\Sigma) \triangleq \text{diag}((\sigma_i - \tau)_+)$, and $(x)_+ = \max(x, 0)$. The solution to (16) is given by singular value thresholding according to the following theorem (Theorem 2.1 in [5])

Theorem 1. *For each $\tau \geq 0$, and $X \in \mathbb{R}^{n_1 \times n_2}$:*

$$D_\tau(X) = \arg\min_{Y \in \mathbb{R}^{n_1 \times n_2}} \left\{ \frac{1}{2} \|Y - X\|_F^2 + \tau \|Y\|_* \right\}. \quad (18)$$

Theorem 1 indicates that the solution to (15) is given by

$$M_k = D_{\lambda/t_k} \left(M_{k-1} - \frac{1}{t_k} \nabla f(M_{k-1}) \right). \quad (19)$$

The remaining question then becomes how to deal with the feasible set Γ_0. Note that (15) is a strongly convex problem, so it is reasonable to project M_k onto the convex set Γ_0 at the kth iteration. For a matrix M, define

$$\mathcal{P}(M) = \frac{I}{\|M\|_{1,1}} M \quad (20)$$

as the projection of M on to the convex set Γ_0. At the kth iteration, we replace M_k obtained from (19) by $\mathcal{P}(M_k)$. Note that at the kth iteration we do not force M_k to be a non-negative matrix and the following initialization explains the reason.

Intuitively, the initialization we choose should be as close as possible to the matrix with maximal likelihood. In other words, we would initialize with a matrix that minimizes $f(M)$ in (10). For this consideration, we initialize by $M_0 = \mathcal{P}(\sum_{i=1}^n y_i A_i)$ (similar to the initialization for alternating minimization in [8]). However, a difference from [9] is that rather than taking the top k singular value, we keep all singular values to preserve information that may be needed for future iterations before truncating them prematurely. In our algorithm, all singular values of M_k decreases as k increases. With such an initialization, the magnitude of the gradient $\nabla[-\log p(y|\mathcal{A}M_k)]$ is typically small at each iteration. Hence, we can ensure each M_k to be nonnegative by choosing a sufficiently small step size at the kth iteration. The algorithm is summarized in Algorithm 1.

Details of Algorithm 1 are as follows. Here L is the step size, $\gamma > 1$ changes the step size to ensure the cost function to decrease

Algorithm 1 PMLSV

1: Initialize: $M_0 = \mathcal{P}(\sum_{i=1}^n y_i A_i)$, parameter γ, step size L
2: **for** $k = 1, 2, \ldots NOI$ **do**
3: $\mathcal{G}(M_{k-1}) := \nabla[-\log p(y|\mathcal{A}M_{k-1})]$
4: $C := M_{k-1} - \frac{1}{L}\mathcal{G}(M_{k-1})$
5: singular value decomposition: $C := UDV^T$
6: $D_{\text{new}} := \text{diag}((\text{diag}(D) - \frac{\lambda}{L})_+)$
7: $W_k := \mathcal{P}(UD_{\text{new}}V^T)$.
8: If $F(M_k) < F(M_{k-1})$, then $k = k + 1$; else $L = \gamma L$, go to 6.
9: If $|F(M_k) - F(M_{k-1})| < 0.5/NOI$, then $k = k - 1$, exit;
10: **end for**

at each iteration, and NOI is the maximum number of iterations. Steps 3-7 generate solution to (15) at the kth iteration. Step 8 examines if the cost function is reduced in the iteration. If the cost function does not decrease, we update the step size by multiplying γ in order to change the singular value more conservatively. In Step 9, if the absolute difference of cost function between consecutive two iterations is less than $0.5/NOI$, then we stop the algorithm. The choice of NOI is user specified: a larger NOI leads to more accurate solution, and a small NOI obtains the solution quickly at the cost of accuracy.

4. EXAMPLES

We use the image of solar flare as example (see [10] for detailed explanation of the data). We break the image into 8 by 8 patches and vectorize each patch to be a column of a new matrix. This new matrix formed by vectorized patches can be well approximated using a low-rank matrix, as demonstrated in Fig. 1. The intensity of the image is $I = 2.37 \times 10^7$. To change SNR of the image, we scale the image intensity by $\alpha \geq 1$. The parameters for the PMLSV algorithm are $L = 10^{-5}$, $\gamma = 1.1$, and $NOI = 2500$.

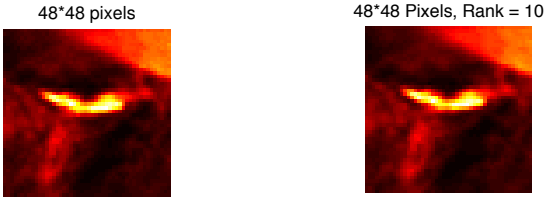

(a) original solar flare image (b) solar flare image with rank 10

Fig. 1: Original and low-rank solar flare image.

First, we run the PMLSV algorithm and solve the SDP using CVX[1] with various number of measurements, respectively, when fixing $\alpha = 4$ and $\lambda = 0.002$. In Fig. 2, the blue line represents the risk by running PMLSV algorithm and red line represents the risk by running SDP given the same observations, respectively. Fig. 2 demonstrates that more measurements lead to more smaller risk, as expected. Also, since it is an approximate algorithm, PMLSV algorithm is less accurate than SDP; however, the maximal increase in risk of PMLSV algorithm relative to that of SDP is 4.89%. PMLSV is much faster: as shown in Table 1 which is the CPU running time of solving SDP by CVX and our PMLSV algorithm.

Second, we run our algorithms with different α when fixing $N = 1000$ and $\lambda = 0.002$. The results are shown in Fig. 3. The

[1] http://cvxr.com/cvx/

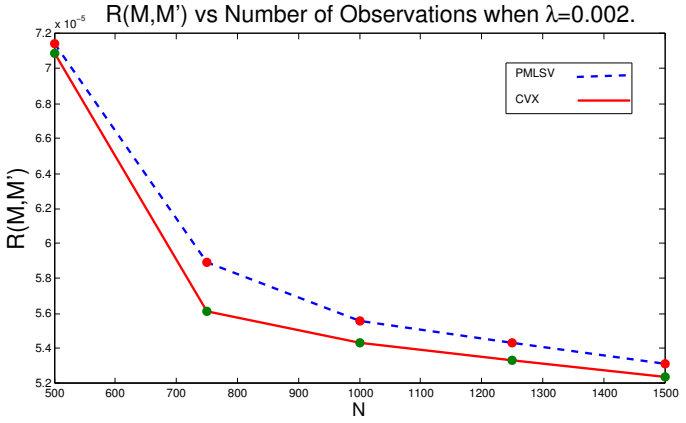

Fig. 2: Risk vs the number of measurements when $\alpha = 4$ and $\lambda = 0.002$. Points from left to right correspond to risk when $N = 500, 750, 1000, 1250, \text{and} 1500$, respectively, and solved by CVX and PMLSV.

Table 1: CPU time (in seconds) of solving SDP by using CVX and our PMLSV algorithm when fixing $\alpha = 4$ and $\lambda = 0.002$ with 500, 750, 1000, 1250 and 1500 measurements, respectively.

N	500	750	1000	1250	1500
SDP	725s	1146s	1510s	2059s	2769s
PMLSV	172s	232s	378s	490s	642s

larger α (hence the higher the SNR) we have, the lower the risk as demonstrated in Fig. 4.

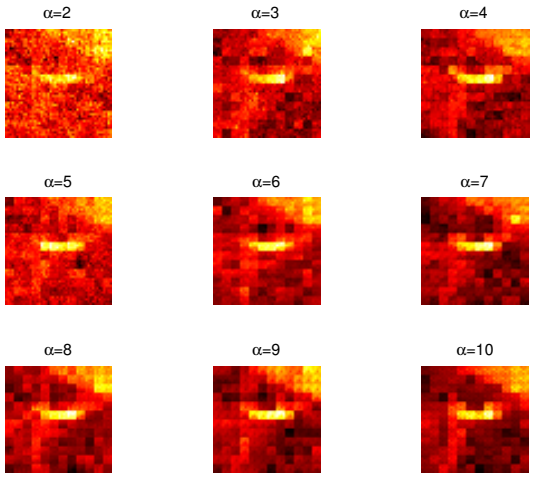

Fig. 3: Recovery results when fixing N =1000, $\lambda = 0.002$ with different value of α.

Third, we run our algorithm with different values of λ when fixing $N = 1000$ and $\alpha = 4$. The results are shown in Fig. 5. From Fig. 6, we can see that there is an optimal value for λ which leads to the smallest risk.

5. CONCLUSION AND FUTURE WORK

We have presented a new algorithm for low-rank matrix recovery with linear measurements contaminated with Poisson noise: the Poisson noise Maximal Likelihood Singular Value Thresholding (PMLSV) algorithm, based on solving a regularized maximum likelihood problem with nuclear norm as the reguarlizer.

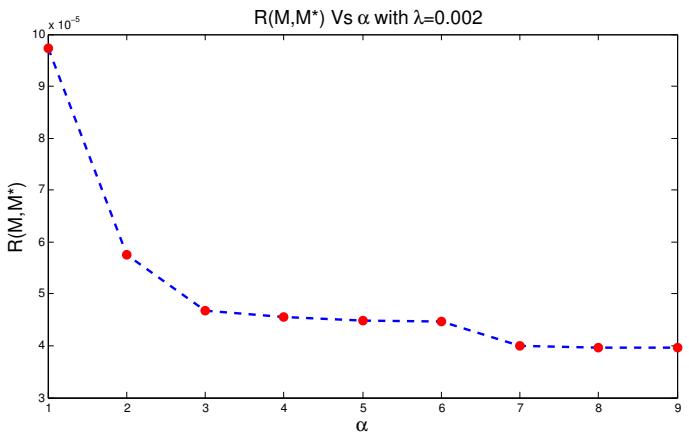

Fig. 4: Risk vs α when fixing $N = 1000$, $\lambda = 0.002$. Points from left to right means the risk with $\alpha = 1$ to $\alpha = 9$.

We demonstrate its accuracy and efficiency compared with the semi-definite program (SDP) and tested on real data examples of solar flare images. Future work include analyzing the convergence property of the algorithm, and extension to the related matrix completion problem with Poisson noise.

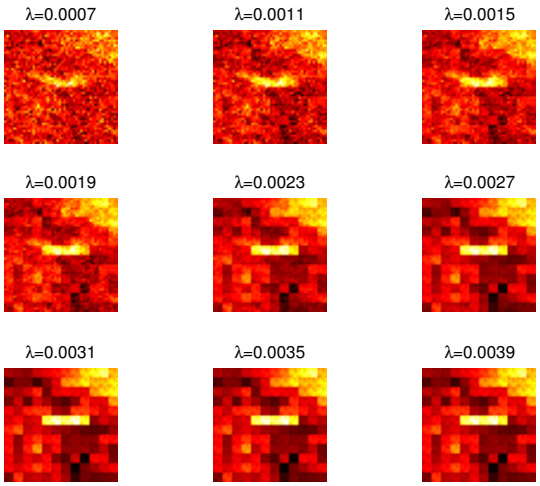

Fig. 5: Recovery results with different value of λ when fixing $N = 1000$ and $\alpha = 4$.

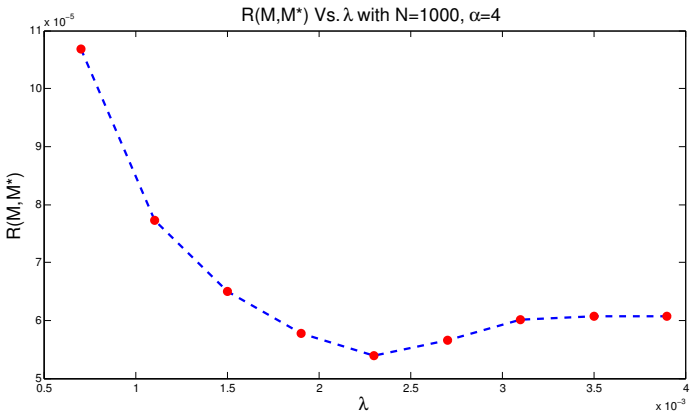

Fig. 6: Risk vs λ when $N = 1000$, $\alpha = 4$. Points from left to right means the risk when λ=0.0007 to λ=0.0039 with step size 0.0004.

6. REFERENCES

[1] David J Brady, *Optical imaging and spectroscopy*, John Wiley & Sons, 2009.

[2] Emmanuel J Candès and Benjamin Recht, "Exact matrix completion via convex optimization", *Foundations of Computational mathematics*, vol. 9, no. 6, pp. 717–772, 2009.

[3] Yao Xie, Yuejie Chi, and Robert Calderbank, "Low-rank matrix recovery with poisson noise", in *IEEE Global Conf. on Signal and Information Processing (GLOBAL SIP)*, June 2013.

[4] Akshay Soni and Jarvis Haupu, "Estimation error guarantees for poisson denoising with sparse and structured dictionary models", in *International Symposium on Information Theory (ISIT)*, 2014.

[5] Jian-Feng Cai, Emmanuel J Candès, and Zuowei Shen, "A singular value thresholding algorithm for matrix completion", *SIAM J. Optimization*, vol. 20, no. 4, pp. 1956–1982, 2010.

[6] Maxim Raginsky, Rebecca M Willett, Zachary T Harmany, and Roummel F Marcia, "Compressed sensing performance bounds under poisson noise", *IEEE Trans. Signal Processing*, vol. 58, no. 8, pp. 3990–4002, 2010.

[7] Yao Xie, Yuejie Chi, Yang Cao, and Robert Calderbank, "Low-rank matrix recovery in poisson noise", *working paper*.

[8] Shuiwang Ji and Jieping Ye, "An accelerated gradient method for trace norm minimization", in *Proc. of 26th Annual Int. Conf. on Machine Learning (ICML)*. ACM, 2009, pp. 457–464.

[9] Prateek Jain, Praneeth Netrapalli, and Sujay Sanghavi, "Low-rank matrix completion using alternating minimization", in *Proc. of 45th ACM Symp. on theory of computing (STOC)*. ACM, 2013, pp. 665–674.

[10] Y. Xie, J. Huang, and R. Willett, "Change-point detection for high-dimensional time series with missing data", *IEEE J. Sel. Top. Signal Processing (J-STSP)*, vol. 7, no. 1, pp. 12–27, Feb. 2013.